

BLACKWELL APPROACHABILITY AND MINIMAX THEORY

MATUS TELGARSKY

ABSTRACT. This manuscript investigates the relationship between Blackwell Approachability, a stochastic vector-valued repeated game, and minimax theory, a single-play scalar-valued scenario. First, it is established in a general setting — one not permitting invocation of minimax theory — that Blackwell’s Approachability Theorem (Blackwell [1]) and its generalization due to Hou [6] are still valid. Second, minimax structure grants a result in the spirit of Blackwell’s weak-approachability conjecture (Blackwell [1]), later resolved by Vieille [11], that any set is either approachable by one player, or avoidable by the opponent. This analysis also reveals a strategy for the opponent.

1. INTRODUCTION.

Consider a repeated game between two players, one selecting $x_t \in \mathcal{X}$, the other $y_t \in \mathcal{Y}$, where the payoff is determined by a vector-valued function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$. In Blackwell’s Approachability Game (Blackwell [1]), the player choosing from \mathcal{X} tries to keep the center of gravity $t^{-1} \sum_{i=1}^t f(x_i, y_i)$ arbitrarily close to some target set $S \subseteq \mathbb{R}^d$, whereas the opponent aims to deny this.

This game has many interesting applications: the existence of Hannan consistent forecasters (Blackwell [2]) and calibrated forecasters (Foster and Vohra [4]), to name a few. But on the abstract side, the basic quandary is: what is the equilibrium structure?

This was a central question of Blackwell’s original manuscript, “An analog of the minimax theorem for vector payoffs” (Blackwell [1]). Both there, and in subsequent presentations, one may find invocations of standard minimax theory; for instance, when establishing Blackwell’s Approachability Theorem, which geometrically characterizes sets where the \mathcal{X} -player can guarantee victory. But the usual setting had linear payoffs and convex compact domains, which are the sufficient conditions for von Neumann’s Minimax Theorem. The question thus remains: what is the real relationship between equilibria in Blackwell Games, and scalar-valued minimax theory?

The goal of this manuscript is to determine this dependence by working in a more general setting (i.e., a choice of $(\mathcal{X}, \mathcal{Y}, f)$ where minimax theorems may simply fail). The organization will be provided shortly, but the main results can be summarized intuitively. Recall that standard minimax theorems — usually providing $\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} f(x, y) = \max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} f(x, y)$ for some family of sets and scalar-valued functions — can be interpreted as stating that the order of the players does not affect the dynamics. This will be carried over to the vector-valued case.

- Hou’s generalization of the Blackwell Approachability Theorem (cf. Hou [6], later rediscovered by Spinat [10]) characterizes the approachable sets when the \mathcal{X} -player moves first; it considers no other order, and thus holds without any appeal to minimax theory. (Cf. Theorem 4.15.)
- When minimax holds for certain scalar-valued subproblems, then the player order may be reversed without changing the dynamics. This will be used to prove a result similar to Vieille’s weak-approachability/weak-excludability theorem (Vieille [11]): Theorem 4.16 will effectively state that minimax structure removes the impact of order in approachability games.

1991 *Mathematics Subject Classification.* Primary: 91A05, 49K35; Secondary: 91A20, 91A15.

Key words and phrases. Blackwell Approachability; vector-valued games; repeated games; minimax theory.

The above characterization made glib reference to scalar-valued subproblems. The most natural appearance of these is to consider halfspaces as target sets, and specifically the scalar-valued game which arises by projecting f onto the halfspace's normal. Such scalarizations have always appeared centrally in the discussion of Blackwell approachability — indeed, Blackwell's approach strategy (which was also used by Hou [6]), is a greedy algorithm that chooses a halfspace orthogonal to the current projection onto the target set, and attempts to force a payoff inside it. A side goal of this manuscript is to provide a deeper understanding of the role played by halfspaces; this turns out to be a convenient detour, as these scalarizations grant an easy mechanism to trace the impact of (scalar-valued) minimax theory.

This manuscript is organized as a progression from minimax problems to the full approachability game. Specifically, after presenting background in Section 2: deterministic single-play vector-valued games are in Section 3; deterministic repeated-play vector-valued games are in Section 4; finally, stochastic repeated-play vector-valued games are in Section 5.

The single-play games of Section 3, though trivial, carry the essential ingredients of the eventual message. First, minimax properties are only needed in the vector-valued case where they are in needed in the scalar-valued cases: precisely when the order of the players must be reversed. Second, this reversal only works for halfspaces: even in the case of compact convex sets, the dynamics become unintelligible.

The repeated games of Section 4 provide the heart of the matter. First, this section develops a geometric characterization of approachable sets in the spirit of Blackwell, Hou, and Spinat [1, 6, 10], but in general spaces which may disallow the application of minimax theorems. Second, in the spirit of weak-approachability/weak-excludability results (Vieille [11]), it provides a characterization of general sets as either approachable by one player, or avoidable by the other. This second result depends on minimax structure, and can fail without it. This section also presents a strategy for the opponent.

To close, Section 5 confirms that studying the deterministic cases suffices: the stochastic and deterministic settings have the same approachable sets.

2. BACKGROUND.

Denote by $\langle \cdot, \cdot \rangle$ and $\| \cdot \|$ the standard inner product and norm on \mathbb{R}^d . For two points $\phi, \psi \in \mathbb{R}^d$, overload interval notation for higher dimensions: e.g., $[\phi, \psi] := \{\alpha\phi + (1 - \alpha)\psi : \alpha \in [0, 1]\}$. Write $\rho(\phi, \psi) := \|\phi - \psi\|$, and for a set $S \subseteq \mathbb{R}^d$, define $\rho(\phi, S) := \inf_{\psi \in S} \rho(\phi, \psi)$. Let $B(\phi, \epsilon)$ denote the closed ball of radius ϵ , and $S_\epsilon := S + B(0, \epsilon)$ be the ϵ -neighborhood of S . Recall that, for nonempty compact sets $S, U \subseteq \mathbb{R}^d$, the Hausdorff metric $\Delta(S, U) := \inf\{\epsilon > 0 : S \subseteq U_\epsilon \wedge U \subseteq S_\epsilon\}$ is complete (cf. Munkres [8, Exercise 45.7]). Finally, for any set S , let S° , \bar{S} , and S^c respectively denote the interior, closure, and complement of S .

This manuscript will consider both single-play vector-valued games, termed *forcing* games, and repeated games, termed approachability games. These games are characterized by a 4-tuple $(\mathcal{X}, \mathcal{Y}, f, S)$.

- One player, the \mathcal{X} -player or simply \mathcal{X} , chooses $x \in \mathcal{X}$; the opponent (similarly \mathcal{Y} -player or just \mathcal{Y}) chooses $y \in \mathcal{Y}$.
- A bounded function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$ with bound $\|f(x, y)\| < \gamma$ determines the payoff.
- The \mathcal{X} -player desires payoffs near a target set $S \subseteq \mathbb{R}^d$; the \mathcal{Y} -player tries to avoid S .

Boundedness of f is necessary to the analysis; note that the strict inequality prevents $\gamma = 0$, which is a trivial scenario anyway.

Definition 2.1. Say $g : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ has the *minimax property* when there exist $\bar{x} \in \mathcal{X}$, $\bar{y} \in \mathcal{Y}$ satisfying

$$\sup_{y \in \mathcal{Y}} g(\bar{x}, y) \leq g(\bar{x}, \bar{y}) \leq \inf_{x \in \mathcal{X}} g(x, \bar{y}).$$

(This provides $\inf_x \sup_y g(x, y) = \sup_y \inf_x g(x, y)$.) A function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$ has the *minimax property* when, for every $\lambda \in \mathbb{R}^d$, the function $\langle f(\cdot, \cdot), \lambda \rangle$ has the minimax property.

This latter property is quite restrictive, and Appendix D discusses function classes which satisfy it. Reliance on the minimax property will always be stated explicitly.

The single-play game is defined as follows; note that, with the exception of Section 5, the games in this manuscript are deterministic.

Definition 2.2. A *forcing game* has only one round, where \mathcal{X} wins iff $f(x, y) \in S$ (\mathcal{Y} wins iff $f(x, y) \in S^c$). Say \mathcal{X} can force S as player 1, or more succinctly \mathcal{X} can 1-force S , when $\exists x \in \mathcal{X} \cdot \forall y \in \mathcal{Y} \cdot f(x, y) \in S$. The weaker property that \mathcal{X} can 2-force S means $\forall y \in \mathcal{Y} \cdot \exists x \in \mathcal{X} \cdot f(x, y) \in S$. Analogously, one can define 1- and 2-forcing for \mathcal{Y} , where the goal is now S^c .

This terminology reflects an ordering of player moves induced by the quantifiers. Suppose \mathcal{X} can 1-force S ; then even if $y \in \mathcal{Y}$ is selected with knowledge of the chosen $\bar{x} \in \mathcal{X}$, an \bar{x} exists so that $f(\bar{x}, y) \in S$. On the other hand, if \mathcal{Y} can 2-force S^c , and if this choice is with knowledge of the selected $x \in \mathcal{X}$, then the outcome $f(x, y) \notin S$ can be forced. If \mathcal{X} can 1-force S , it can win with either order. But if \mathcal{X} can only 2-force S , then in general it can only win as the player who moves last.

In the repeated game setting, players choose $(x_t)_{t=1}^\infty$ and $(y_t)_{t=1}^\infty$, and have access to the selection history $\mathcal{H}_t := ((x_i, y_i))_{i=1}^t$. A strategy for \mathcal{X} is a family of functions $\mathbf{g} := (\mathbf{g}_t : (\mathcal{X} \times \mathcal{Y})^{t-1} \rightarrow \mathcal{X})_{t=1}^\infty$; analogously, a strategy for \mathcal{Y} is a family $\mathbf{h} := (\mathbf{h}_t : (\mathcal{X} \times \mathcal{Y})^{t-1} \rightarrow \mathcal{Y})_{t=1}^\infty$. In round $t+1$, the players are presented the current history $\mathcal{H}_t := ((x_i, y_i))_{i=1}^t$, then select $x_{t+1} := \mathbf{g}_{t+1}(\mathcal{H}_t)$ and $y_{t+1} := \mathbf{h}_{t+1}(\mathcal{H}_t)$.

Given a history \mathcal{H}_t , define $\phi_t := t^{-1} \sum_{i=1}^t f(x_i, y_i)$. the goal in the (deterministic) repeated setting will be to manipulate this center of gravity, where the averaging will suffice to make the family of approachable sets much richer than the forcible sets.

Definition 2.3. A set S (or, for clarity, a tuple $(\mathcal{X}, \mathcal{Y}, f, S)$) is *approachable* when

$$\exists \mathbf{g} \cdot \forall \epsilon > 0 \cdot \exists T \cdot \forall \mathbf{h} \cdot \forall t \geq T \cdot \phi_t \in S_\epsilon.$$

Any \mathbf{g} satisfying this definition is an *approach strategy*.

In other words, $\lim_{t \rightarrow \infty} \inf_{z \in S} \|\phi_t - z\| = 0$, and this convergence is uniform with respect to the family of opponent strategies.

Note that for any \mathcal{X} -player strategy \mathbf{g} , there exists a family of opponent strategies which assume that \mathcal{X} is playing \mathbf{g} ; these strategies effectively choose y_t knowing x_t . Moreover, these strategies are considered in the universal quantification over opponent strategies in the definition of approachability. As such, when constructing an approach strategy, what the \mathcal{X} -player can force in each round are the 1-forcible sets; analogously, the opponent strategy is working with 2-forcible sets. Said another way, the quantification order in the definition of approachability implies another setting where \mathcal{X} moves first, and \mathcal{Y} observes this before moving.

It is thus natural to consider the effect of reordering these quantifiers. First, one can say S is not approachable when

$$\forall \mathbf{g} \cdot \exists \epsilon > 0 \cdot \forall T \cdot \exists \mathbf{h} \cdot \exists t \geq T \cdot \phi_t \notin S_\epsilon.$$

Reversing the quantifiers for \mathbf{g} and \mathbf{h} yields a setting where \mathcal{Y} moves first.

Definition 2.4. A set S (or, for clarity, a tuple $(\mathcal{X}, \mathcal{Y}, f, S)$) is *avoidable* when

$$\exists \mathbf{h} \cdot \exists \epsilon > 0 \cdot \forall T \cdot \forall \mathbf{g} \cdot \exists t \geq T \cdot \phi_t \notin S_\epsilon.$$

Any \mathbf{h} satisfying this definition is an *avoidance strategy*.

Mirroring the discussion of player order and approachability, an avoidance game effectively has the \mathcal{Y} player move first, and \mathcal{X} chooses with knowledge of this move. Correspondingly, in any given round, \mathcal{Y} can force 1-forcible sets, and \mathcal{X} has the easier criterion of 2-forcible sets.

Blackwell considered a stronger property than avoidability, called *excludability*, where the quantifiers $\exists \epsilon > 0$, $\forall T$, and $\exists t \geq T$ were respectively replaced with $\forall \epsilon > 0$, $\exists T$, and $\forall t \geq T$, thus matching the goal of the \mathcal{X} -player. While there exist games which are neither approachable nor excludable (Blackwell [1]), a weaker definition grants that every game is either weak-approachable

or weak-excludable (Vieille [11]). Section 4 will show that, under the minimax property, every game is either approachable or avoidable; as in the scalar-valued setting, minimax structure nullifies the effect of player order.

A few final technical conveniences are in order. Since S is approachable iff its closure is approachable, this manuscript will follow the usual convention of considering only the case that S is closed. Next, every superset S' of an approachable set $S \subseteq S'$ is itself approachable, simply by running the approach strategy for S . Combining this with the fact that f is bounded, it suffices to consider only compact sets. Lastly, define $\text{conv}(f(\mathcal{X}, \mathcal{Y}))$ to be the convex hull of the range of f ; notice that $\phi_t \in \text{conv}(f(\mathcal{X}, \mathcal{Y}))$.

3. FORCING GAMES.

A few properties are straight from the definitions.

- \mathcal{X} can 1-force S iff \mathcal{Y} can not 2-force S^c .
- If \mathcal{X} can 1-force S , then it can 1-force $S' \supseteq S$.
- \mathcal{X} can 1-force S iff S intersects every set S' which can be 2-forced by \mathcal{Y} .
- If \mathcal{X} can 1-force S , then \mathcal{X} can 2-force S .

Attempting to reverse this final property is where things become interesting.

Proposition 3.1. *Let any halfspace $H := \{z \in \mathbb{R}^d : \langle \lambda, z \rangle \leq c\}$ be given, and suppose $\langle f(\cdot, \cdot), \lambda \rangle$ has the minimax property. If \mathcal{X} can 2-force H , then \mathcal{X} can 1-force H .*

Proof. Given any $y \in \mathcal{Y}$, choose $x_y \in \mathcal{X}$ satisfying $f(x_y, y) \in H$. Since for all y

$$c \geq \langle f(x_y, y), \lambda \rangle \geq \inf_{x \in \mathcal{X}} \langle f(x, y), \lambda \rangle,$$

it follows that

$$c \geq \sup_{y \in \mathcal{Y}} \inf_{x \in \mathcal{X}} \langle f(x, y), \lambda \rangle = \sup_{y \in \mathcal{Y}} \langle f(\bar{x}, y), \lambda \rangle,$$

where the existence of \bar{x} is from the minimax property. It follows that for every $y \in \mathcal{Y}$, $f(\bar{x}, y) \in H$. \square

To see that minimax structure is necessary, it suffices to consider a game in \mathbb{R}^1 with $a := \inf_x \sup_y f(x, y) > \sup_y \inf_x f(x, y) =: b$; in particular, the set $H := (-\infty, (a+b)/2]$ can be 2-forced by \mathcal{X} , but not 1-forced. This problem can then be lifted to any dimension $d > 1$ by constructing $f' : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$ with $\langle f'(\cdot, \cdot), \lambda \rangle = f(\cdot, \cdot)$, and considering the set $H' := \{z \in \mathbb{R}^d : \langle \lambda, z \rangle \leq (a+b)/2\}$. This final example also demonstrates how the restriction to halfspaces fits in: the game is effectively projected onto the halfspace normal, thus becoming scalar-valued.

So the natural question is: what can be said about sets which are not halfspaces?

3.1. Vector-valued Games as 3-player Games. Even in the case of minimax structure and compact convex target sets, there are 2-forcible sets which can not be 1-forced: such an example appears in Appendix A. The goal of this subsection is to investigate why these difficulties arise.

First note the following characterization of (closed convex) set membership.

Proposition 3.2. *Let a closed convex nonempty set S and any point $\phi \in \mathbb{R}^d$ be given. Let $\sigma_S(\lambda) = \sup_{z \in S} \langle z, \lambda \rangle$ denote the support function of S . Then*

$$\phi \in S \iff \sup_{\substack{\lambda \in \mathbb{R}^d \\ \|\lambda\| \leq 1}} \langle \phi, \lambda \rangle - \sigma_S(\lambda) = 0.$$

As $\phi \in S$ is equivalent to $\inf_{z \in S} \|\phi - z\|^2 = 0$, this statement can be understood via convex duality.




FIGURE 1. Depiction of Definition 4.1.

Proof. In general, $\sup_{\|\lambda\| \leq 1} \langle \phi, \lambda \rangle - \sigma_S(\lambda) \geq \langle \phi, 0 \rangle - \sigma_S(0) = 0$. Suppose $\phi \in S$; then for any λ , $\langle \phi, \lambda \rangle \leq \sup_{\phi \in S} \langle \phi, \lambda \rangle = \sigma_S(\lambda)$, thus $\sup_{\lambda} \langle \phi, \lambda \rangle - \sigma_S(\lambda) \leq 0$. On the other hand, if $\phi \notin S$, there exists a halfspace $H = \{z \in \mathbb{R}^d : \langle \lambda', z \rangle \leq c\}$ satisfying $S \subseteq H$ and $\phi \notin H$, meaning $\sigma_S(\lambda') \leq \sigma_H(\lambda') = c < \langle \phi, \lambda' \rangle$ so $\langle \phi, \lambda' \rangle - \sigma_S(\lambda') > 0$. \square

For any closed convex set S , consider the scalar valued function

$$g_S(x, y, \lambda) := \langle f(x, y), \lambda \rangle - \sigma_S(\lambda).$$

Set $\mathcal{Z} := \{\lambda \in \mathbb{R}^d : \|\lambda\| \leq 1\}$; Proposition 3.2 grants

$$\begin{aligned} \mathcal{X} \text{ can 1-force } S &\implies \inf_{x \in \mathcal{X}} \sup_{y \in \mathcal{Y}} \sup_{\lambda \in \mathcal{Z}} g_S(x, y, \lambda) = 0, \\ \mathcal{X} \text{ can 2-force } S &\implies \sup_{y \in \mathcal{Y}} \inf_{x \in \mathcal{X}} \sup_{\lambda \in \mathcal{Z}} g_S(x, y, \lambda) = 0. \end{aligned}$$

In general, $\inf_x \sup_y \sup_{\lambda} f(x, y, \lambda) \geq \sup_y \inf_x \sup_{\lambda} f(x, y, \lambda)$. But when this does not hold with equality for a 2-forcible S , one has an example which is 2-forcible but not 1-forcible. The remaining discussion thus considers this expression.

Suppose now that \mathcal{X}, \mathcal{Y} are convex and compact, and that $f(\cdot, y)$ is convex for every $y \in \mathcal{Y}$, and $f(x, \cdot)$ is concave for every $x \in \mathcal{X}$. Since $\langle \cdot, \cdot \rangle$ is bilinear and σ_S is a convex function (cf. Hiriart-Urruty and Lemaréchal [5, Proposition B.2.1.2]), g_S is concave in λ .

Indeed, g_S can be viewed as providing the payoff for a 3-player game. The aforementioned structure suffices to grant one type of equilibrium: in particular, the existence of Nash Equilibria (cf. Borwein and Lewis [3, Exercise 8.3.10.d] for a suitable generalization). But that only grants that players have no incentive to deviate without collusion, whereas here y and λ are effectively cooperating. The function $\sup_{\lambda} g_S(x, \cdot, \lambda)$ is highly nonconvex, and thus lacks the usual structure allowing a statement like $\inf_x \sup_y \sup_{\lambda} g_S(x, y, \lambda) = \sup_y \inf_x \sup_{\lambda} g_S(x, y, \lambda)$. Thus, in general, there is a gap between the two sides of this expression.

This third player choosing λ thus introduces major difficulty into the problem. Note that in the repeated game, only one λ is chosen (not a sequence, as with x_t and y_t). This suffices to make the approachable sets far different from the forcible sets. Furthermore, the strategies for both players can be seen as attempting to work with or against this third player λ , whose maximizing choice is the direction of projection onto S , equivalently a hyperplane tangent to S .

4. APPROACHABILITY GAMES.

Halfspaces continue to be central in the repeated game setting; the following definition provides the basic tool whereby a halfspace is useful to either player.

Definition 4.1. Consider, as in Figure 1, a pair of points $\psi \in S$ and $\phi \in \text{conv}(f(\mathcal{X}, \mathcal{Y})) \setminus S$ with $\rho(\phi, \psi) = \rho(\phi, S)$, and a halfspace H orthogonal to and also passing through $[\phi, \psi]$. Refer to any tuple (ϕ, ψ, H) satisfying this arrangement as a *halfspace-forcing candidate for S* . If \mathcal{X} can 1-force H , call this a *halfspace-forcing example for S* . When \mathcal{X} can not 1-force H (meaning \mathcal{Y} can 2-force H^c), call this tuple a *halfspace-forcing counterexample for S* .

Consider a halfspace-forcing candidate (ϕ, ψ, H) for some compact set S . If this is a halfspace-forcing example, and $\phi = \phi_t = t^{-1} \sum_{i=1}^t f(x_i, y_i)$ for some iteration t , then the \mathcal{X} player may force H and move closer to S ; this will be proved in Lemma 4.3, and will provide the basis for the approach strategy \mathbf{g}^* . On the other hand, if this is a halfspace-forcing counterexample, and ϕ_t is close to ψ for some iteration t , then the opponent may force H^c and move away from S ; this is proved in Lemma 4.8, and provides the basis for the opponent strategy \mathbf{h}^* .

Halfspace-forcing examples were used by Blackwell [1] to construct the original greedy approach strategy (and were again used by Hou [6] and Spinat [10]); the only distinction here is that the halfspaces need not touch the target S . Where the present work truly departs from earlier works is by also developing a theory of unapproachable sets with halfspaces as a starting point. This will allow the construction of an opponent strategy \mathbf{h}^* , and also make it easy to measure the impact of minimax structure (since, combined with Proposition 3.1 and properties of forcible sets, if (ϕ, ψ, H) is a halfspace-forcing counterexample, then \mathcal{Y} may equivalently 1-force and 2-force H^c), culminating in the approachability/avoidability statement of Theorem 4.16. The following definition, which extends halfspace-forcing into a global property, will characterize the approachable sets in Theorem 4.15.

Definition 4.2. A set $S \subseteq \mathbb{R}^d$ is an *A-set* if it has no halfspace-forcing counterexamples.

In the case that \mathcal{X} is compact and $f(\cdot, y)$ is continuous for every $y \in \mathcal{Y}$, an A-set S will also satisfy a stronger property that halfspace-forcing candidates (ϕ, ψ, H) may place H tangent to S . Indeed, this grants a type of set called a *B-set* by Spinat [10], which was used there to characterize approachable sets. In that setting, however, the compactness and continuity were guaranteed; in the more general choice of $(\mathcal{X}, \mathcal{Y}, f)$ here, the relaxed notion of A-set is necessary and sufficient.

4.1. Sufficient Conditions for Approachability. First, a quantification of the progress granted by a single halfspace-forcing example.

Lemma 4.3. *Let (ϕ, ψ, H) be a halfspace-forcing example for S , and set $\tau := \rho(\psi, H^c)$. Then there exists $\bar{x} \in \mathcal{X}$ so that, for any $y \in \mathcal{Y}$,*

$$\langle f(\bar{x}, y) - \psi, \phi - \psi \rangle \leq \tau\gamma.$$

Proof. Set ψ' to be the projection of ϕ onto H , and choose any $\bar{x} \in \mathcal{X}$ providing 1-forcibility of H . Then, for any $y \in \mathcal{Y}$,

$$\langle f(\bar{x}, y) - \psi, \phi - \psi \rangle = \langle f(\bar{x}, y) - \psi', \phi - \psi \rangle + \langle \psi' - \psi, \phi - \psi \rangle \leq 0 + \tau\gamma,$$

which made use of $\|\phi - \psi\| \leq \gamma$. \square

This lemma leads to the following greedy strategy $\mathbf{g}^* = (\mathbf{g}_t^*)_{t=1}^\infty$ for the approach player, parameterized by a family of tolerances $(\tau_t)_{t=1}^\infty$ (these tolerances being the only modification to the strategy provided by Blackwell, Hou, and Spinat [1, 6, 10]).

Definition 4.4. The \mathcal{X} -player strategy $\mathbf{g}^* = (\mathbf{g}_t^*)_{t \geq 1}$ is

$$\mathbf{g}_{t+1}^*(\mathcal{H}_t) := \begin{cases} x & \exists x, \psi_t, H_t \bullet (\phi_t, \psi_t, H_t) \text{ is a halfspace-forcing example for } S, \\ & \rho(\psi_t, H_t^c) \leq \tau_t, \forall y \in \mathcal{Y} \bullet f(x, y) \in H; \\ \text{arbitrary} & \text{otherwise.} \end{cases}$$

Theorem 4.5. *Let an A-set S , any tolerances $(\tau_t)_{t=1}^\infty$ with $\sum_{i \geq 1} \tau_i \leq \gamma$ and $\tau_t > 0$, and any $\epsilon > 0$ be given. If \mathcal{X} uses strategy \mathbf{g}^* , then for any opponent strategy and any $t \geq 3\gamma^2/\epsilon^2$, $\phi_t \in S_\epsilon$.*

The proof reveals that $\sum_{i=1}^t \tau_i = o(t)$ suffices; requiring a constant bound is for convenience.

Proof. So that ψ_t is always defined, set $\psi_t := \phi_t$ when $\phi_t \in S$. It will be shown by induction that

$$\|\phi_t - \psi_t\|_2^2 \leq \frac{\gamma^2 + 2\gamma \sum_{i=1}^{t-1} \tau_i}{t};$$




FIGURE 2. Depiction of Lemma 4.8.

the result follows by applying the bounds for $\sum_i \tau_i$ and t , then taking a square root.

In the base case, $\|\phi_1 - \psi_1\| \leq \gamma$. For the inductive step,

$$\begin{aligned}
 \|\phi_{t+1} - \psi_{t+1}\|_2^2 &\leq \|\phi_{t+1} - \psi_t\|_2^2 \\
 &= \frac{1}{(t+1)^2} \|t(\phi_t - \psi_t) + (f(x_{t+1}, y_{t+1}) - \psi_t)\|_2^2 \\
 (4.6) \quad &= \frac{1}{(t+1)^2} (t^2 \|\phi_t - \psi_t\|_2^2 + 2t \langle \phi_t - \psi_t, f(x_{t+1}, y_{t+1}) - \psi_t \rangle + \|f(x_{t+1}, y_{t+1}) - \psi_t\|_2^2).
 \end{aligned}$$

If $\phi_t \in S$, then $\phi_t - \psi_t = 0$ and the middle term vanishes. Otherwise, since $\tau_t > 0$ and S is an A-set, there must exist a halfspace-forcing example (ϕ_t, ψ_t, H_t) with $\rho(\psi_t, H_t^c) \leq \tau_t$, and Lemma 4.3 grants $\langle \phi_t - \psi_t, f(x_{t+1}, y_{t+1}) - \psi_t \rangle \leq \tau_t \gamma$. Applying the inductive hypothesis,

$$(4.6) \leq \frac{t\gamma^2 + 2\gamma t \sum_{i=1}^t \tau_i + \gamma^2}{(t+1)^2} \leq \frac{\gamma^2 + 2\gamma \sum_{i=1}^t \tau_i}{t+1}.$$

□

Remark 4.7. This statement also grants the approachability of any set S' which contains an A-set S : simply run \mathbf{g}^* on S' . But it is unclear how to produce an approach strategy for S' directly. Suppose S' is the union of an A-set, and another set which is nearly an A-set: a small piece is missing in such a way to render it unapproachable. The approach strategy must somehow rule out gravitating towards the second, damaged set; detecting this difficulty does not appear tractable. Please see the examples of Appendix A for further discussion.

4.2. Sufficient Conditions for Non-approachability. The first step is complementary to Lemma 4.3: quantifying how much a single halfspace-forcing counterexample benefits the \mathcal{Y} -player.

Lemma 4.8. *Let (ϕ, ψ, H) be a halfspace-forcing counterexample. Set $\tau := \rho(\psi, H^c) > 0$, $\epsilon := \frac{\tau^2(\sqrt{\gamma^2 + \tau^2} - \gamma)}{8(4\gamma^2 + \tau^2)}$ and let $T \geq \lceil 8/\epsilon \rceil$ be given. Then for any $p \in B(\psi, 2\epsilon)$, and any sequence $(x_i)_{i \geq 1}$, there exist $(y_i)_{i \geq 1}$ and $M \leq \lceil T\gamma\epsilon/8 \rceil$ where*

$$\rho\left(\phi, (T+M)^{-1}(Tp + \sum_{i=1}^M f(x_i, y_i))\right) \leq \rho(\phi, S) - \epsilon.$$

Under the stronger condition that H^c may be 1-forced by \mathcal{Y} , there exists a single \bar{y} so that the choice $y_i = \bar{y}$ grants the above property for any sequence $(x_i)_{i \geq 1}$.

This derivation is mechanical, and thus pushed to Appendix B.1, but the idea, which appears in Figure 2, is simple. By assumption, there exists a point $\psi \in S$ satisfying $\psi \in B(\phi, \rho(\phi, S)) \cap S$, and a halfspace H whose boundary is τ away from ψ , and H is not 1-forcible by \mathcal{X} . Correspondingly, H^c is 2-forcible by \mathcal{Y} .

To see how this helps \mathcal{Y} , by properties of l^2 balls and boundedness of f , every line connecting ψ to $H^c \cap f(\mathcal{X}, \mathcal{Y})$ must pass interior to $B(\phi, \rho(\phi, S))$. This remains true for a tiny neighborhood

around ψ . Thus it suffices for the \mathcal{V} player to force points in H^c : regardless of how the \mathcal{X} player chooses, a future center of gravity will eventually land interior to $B(\phi, \rho(\phi, S))$. The extra work in the lemma is in producing ϵ as a function of τ , allowing uniform controls for various halfspace-forcing counterexamples.

Definition 4.9. Define $\epsilon(\tau)$ as in the statement of Lemma 4.8. For any set $S \subseteq \mathbb{R}^d$, define the *excess of S with tolerance $\tau > 0$* as

$$\mathcal{E}_\tau(S) := \left\{ B(\psi, \epsilon(\tau))^o : \exists \phi, \psi, H. (\phi, \psi, H) \text{ is a halfspace-forcing counterexample with } \rho(\psi, H^c) \geq \tau \right\}.$$

Define $\mathcal{V}_\tau(S) := S \setminus \bigcup_{U \in \mathcal{E}_\tau(S)} U$.

As per the following lemma, the operator \mathcal{V}_τ removes points which can not be used by an approach strategy. A similar tool and subsequent limiting argument appeared in the analysis of Hou [6] and Spinat [10], albeit without τ , which will be used in constructing the opponent strategy.

Lemma 4.10. *For any $\tau > 0$ and $S \subseteq \mathbb{R}^d$, S is approachable iff $\mathcal{V}_\tau(S)$ is approachable.*

Proof. (\implies) Suppose S is approachable, but contradictorily that $S' := \mathcal{V}_\tau(S)$ is not approachable. Necessarily, $\mathcal{V}_\tau(S) \neq S$ and $\mathcal{E}_\tau(S)$ is nonempty. Let \mathbf{g} be any approach strategy for S ; it must fail to approach S' , and therefore there must exist an opponent strategy \mathbf{h} such that $\phi_t \in S \setminus S'$ for infinitely many t . Now choose $T_0 \geq \lceil 8/\epsilon(\tau) \rceil$, and let T_1 be the value provided by \mathbf{g} guaranteeing $\phi_t \in S_{\epsilon(\tau)}$ for all opponent strategies and $t \geq T_1$. Finally, set $T := \max\{T_0, T_1\}$. Consider the modified strategy \mathbf{h}' which executes as \mathbf{h} until some $t > T$ satisfies $\phi_t \in S \setminus S'$. Thereafter, it follows the choices granted by Lemma 4.8 (with $p = \phi_t$), and thus guarantees the existence of $t' \geq T$ with $\phi_{t'} \notin S_{\epsilon(\tau)}$, contradicting the approachability of S by \mathbf{g} . But \mathbf{g} was arbitrary, and thus S is not approachable.

(\impliedby) Supersets of approachable sets are always approachable. \square

Theorem 4.11. *Let compact $S \subseteq \mathbb{R}^d$ be given, set $S_0 := S$ and $S_{i+1} := \mathcal{V}_{1/(i+1)}(S_i)$. Then the limit S_∞ (in Hausdorff metric) exists. Furthermore, exactly one of the following statements holds:*

- (1) *There exists $N \in \mathbb{N}$ with $S_n = \emptyset$ for all $n \geq N$, and each S_i is not approachable;*
- (2) *S_∞ is a compact nonempty A-set.*

The tolerance $1/(i+1)$ will make it easy to control the behavior of the eventual opponent strategy \mathbf{h}^* .

Proof. Suppose there exists some N such that $S_N = \emptyset$; since \mathcal{V} only makes sets smaller, it follows that $S_n = \emptyset$ when $n \geq N$, providing $S_\infty = \emptyset$. Next, the empty set is never approachable, thus S_n is not approachable when $n \geq N$. But Lemma 4.10 grants that S_i is approachable iff S_{i+1} is approachable, and so by N applications of this lemma, it follows that every S_i is not approachable.

Now suppose there does not exist any N with $S_N = \emptyset$. Thus each S_i is compact and nonempty, and S_∞ exists and is a compact nonempty set by the completeness of the Hausdorff metric on compact nonempty sets.

Assume contradictorily that S_∞ is not an A-set, meaning it has a halfspace-forcing counterexample (ϕ, ψ, H) , and set $\tau := \rho(\psi, H^c) > 0$. By Proposition B.6, there exists $\delta > 0$ so that any set S' satisfying $\Delta(S', S_\infty) < \delta$ has a halfspace-forcing counterexample (ϕ', ψ', H') with $\rho(\psi', (H')^c) \geq \tau/4$.

Choose j so that $\Delta(S_\infty, S_j) < \min\{\delta, \epsilon(\tau/4)/2\}$. Thus $\mathcal{E}_{1/(j+1)}(S_j)$ contains $B(\psi', \epsilon(\tau/4))^o$. But \mathcal{V} only shrinks sets, meaning $\Delta(S_\infty, S_j) \geq \Delta(S_{j+1}, S_j) \geq \epsilon(\tau/4)$, a contradiction. \square

Consider the case that S is not approachable; the above theorem grants the sequence $(S_i)_{i=0}^N$ with $S_0 = S$ and $S_N = \emptyset$. Now define $E_i := S_i \setminus S_{i+1}$ for $i \in (0, \dots, N-1)$, and note that each E_i is disjoint, and $S = \bigcup_{i=1}^{N-1} E_i$. In this way, the operator \mathcal{V} peels S into a finite sequence of concentric sets, like the rinds of an onion; the strategy \mathbf{h}^* , depicted in Figure 3, is to use Lemma 4.8 to move through these rinds, eventually exiting S entirely. It is interesting to contrast the complexity of \mathbf{h}^* with the triviality of \mathbf{g}^* .




FIGURE 3. \mathcal{Y} -player strategy \mathfrak{h}^* , escaping S rind by rind. The small circles denote the regions surrounding specific halfspace-forcing counterexamples, within which an iterate may fall and still make use of the counterexample (cf. Lemma 4.8). These balls are depicted as exceeding the onion rinds since Lemma 4.8, which determines $\delta(S)$, only provides a lower bound on their size. The arrow, abstractly indicating the path out of S , may consist of a tremendous number of iterations.

Definition 4.12. Fix any compact set S , and let $(S_i)_{i \geq 0}$ be as in Theorem 4.11. When there exists N such that $S_N = \emptyset$, define $\delta(S) := \epsilon(1/N)$ and $T(S) := \lceil 8/\delta(S) \rceil$; otherwise, $\delta(S) := 0$ and $T(S) := \infty$. Finally, for any ϕ ,

$$\mathcal{I}_S(\phi) := \begin{cases} -1 & \text{when } \delta(S) > 0 \text{ and } \phi \notin S_{\delta(S)}, \\ \max\{i : \phi \in (S_i)_{\delta(S)}\} & \text{when } \delta(S) > 0 \text{ and } \phi \in S_{\delta(S)}, \\ \infty & \text{otherwise.} \end{cases}$$

Definition 4.13. The \mathcal{Y} -player strategy $\mathfrak{h}^* = (\mathfrak{h}_t^*)_{t \geq 1}$ is

$$\mathfrak{h}_{t+1}^*(\mathcal{H}_t) := \begin{cases} \text{arbitrary} & \text{when } t < T(S); \\ y & \begin{aligned} &\text{when } t = T(S), \text{ or } t > T(S) \text{ and } \mathcal{I}_S(\phi_t) \neq \mathcal{I}_S(\phi_{t-1}), \\ &\quad \text{there exists a halfspace-forcing counterexample } (\phi, \psi, H) \\ &\quad \text{with } \rho(\phi_t, \psi) \leq 2\epsilon(1/(\mathcal{I}_S(\phi_t) + 1)) \text{ and } \rho(\psi, H^c) \geq 1/(\mathcal{I}_S(\phi_{t-1}) + 1), \\ &\quad \text{choose any } y \text{ given by Lemma 4.8 (with } p = \phi_t), \\ &\quad \text{satisfying 1-forcing by } \mathcal{Y} \text{ if possible;} \end{aligned} \\ y' & \begin{aligned} &\text{when } t > T(S) \text{ and } \mathcal{I}_S(\phi_t) = \mathcal{I}_S(\phi_{t-1}), \\ &\quad \text{choose } (\phi, \psi, H) \text{ and } p \text{ as for } \phi_{t-1}, \\ &\quad \text{choose } y' \text{ according to Lemma 4.8, again trying to 1-force.} \end{aligned} \end{cases}$$

A similar mechanism was used by Hou [6, proof of Theorem 3] when proving necessary and sufficient conditions for approachability (that is, the nonconvex form of Blackwell's Approachability Theorem, presented here as Theorem 4.15). There, halfspaces were not used as a primitive tool for the opponent: rather, the strategy was built up by considering a more abstract set guaranteeing progress to the opponent (cf. Hou's *insufficient subsets* [6, Definition 1]). Using halfspaces grants an arguably constructive opponent strategy, but more importantly provides the backbone for measuring the effect of minimax structure.

Proposition 4.14. Let $S \subseteq \mathbb{R}^d$ be given, and suppose $T(S) < \infty$. Then for any approach strategy and any T , there exists $t \geq T$ so that, when $(y_i)_{i \geq 1}$ is chosen by \mathfrak{h}^* ,

$$\phi_t \notin S_{\delta(S)}.$$

Moreover, if every halfspace-forcing counterexample (ϕ, ψ, H) encountered by \mathfrak{h}^* is 1-forcible by \mathcal{Y} , then \mathfrak{h}^* is an avoidance strategy.

Proof. From $T(S) < \infty$ it follows that $\delta(S) > 0$, and for any $\phi \in S_{\delta(S)}$, $\mathcal{I}_S(\phi) < \infty$. Suppose contradictorily that there exists $T \geq T(S)$ so that $\phi_t \in S_{\delta(S)}$ for all $t \geq T$. In particular, this means $\min_{t \geq T} \mathcal{I}_S(\phi_t) =: j > -1$.

Let $t \geq T(S)$ be the earliest iteration with $\mathcal{I}_S(\phi_t) = j$, which by definition of \mathcal{I}_S and the fact $S_i \supseteq S_{i+1}$ grants $\phi_t \in (S_j)_{\delta(S)} \setminus (S_{j+1})_{\delta(S)}$. This ϕ_t satisfies the middle condition in the definition of \mathfrak{h}^* ; it must be verified that the conditions for Lemma 4.8, namely the existence of a halfspace-forcing counterexample and nearby point p , are actually satisfied.

There are two cases for the location of ϕ_t : either it is in $\mathcal{E}_{1/(j+1)}(S_j)$, or it is in $(\mathcal{E}_{1/(j+1)}(S_j))_{\epsilon(S)}$. Recall that $\mathcal{E}_{1/(j+1)}(S_j)$ is the union of balls of radius at least $\epsilon(1/(j+1)) \geq \epsilon(1/N) = \delta(S)$, where the center ψ of each can be made into a halfspace-forcing counterexample (ϕ, ψ, H) . It thus suffices to take $p = \phi_t$: when $\phi_t \in \mathcal{E}_{1/(j+1)}(S_j)$, just take some ball in $\mathcal{E}_{1/(j+1)}(S_j)$ containing ϕ_t , otherwise choose the closest ball, and the triangle inequality gives what is needed (since the conditions on Lemma 4.8 allow distances up to $2\epsilon(1/(j+1))$, and $\delta(S) \leq \epsilon(1/(j+1))$).

As such, $\mathcal{I}_S(\phi_{t+1}) \leq \mathcal{I}_S(\phi_t)$. It is a contradiction if the inequality is strict, so treat it as an equality. But this will cause a chain of iterations all landing in the final case of the definition of \mathfrak{h}^* . Since the same halfspace-forcing counterexample is used, this chain fits exactly with the conditions of Lemma 4.8, and thus some eventual iteration $t' > t$ will satisfy $\mathcal{I}_S(\phi_{t'}) < \mathcal{I}_S(\phi_t)$, a contradiction. Since t was the least counterexample, there are no counterexamples, and the result follows.

Finally, suppose the extra condition that every halfspace-forcing counterexample (ϕ, ψ, H) has H^c 1-forcible by \mathcal{V} . Then the stronger guarantee of Lemma 4.8 is at play, and $\phi_t \notin S_{\delta(S)}$ can be guaranteed regardless of the choice of $(x_t)_{t \geq 1}$. \square

4.3. Approachability and Minimax Theory. Combining the results so far grants a proof of Blackwell's Approachability Theorem in the general (nonconvex) form of Hou [6] and Spinat [10], but in the general setting of this manuscript, with no reliance on minimax structure.

Theorem 4.15. *Let $S \subseteq \mathbb{R}^d$ be given. Then S is approachable iff S contains a compact A-set.*

Proof. Without loss of generality, S may be presumed compact. If S contains an A-set, then Theorem 4.5 provides approachability. If S does not contain an A-set, then $S_\infty \subseteq S$ is not an A-set, and Theorem 4.11 provides non-approachability. \square

The next question then, is what is gained by minimax structure? The answer goes back to Proposition 3.1: minimax structure allows implies that 2-forcible halfspaces are also 1-forcible. This effect propagates to approach games: minimax structure makes the player order inconsequential.

Theorem 4.16. *Suppose $(\mathcal{X}, \mathcal{V}, f)$ has minimax structure, and let $S \subseteq \mathbb{R}^d$ be given. Exactly one of the following statements holds:*

- (1) S is approachable with strategy \mathfrak{g}^* ;
- (2) S is avoidable with strategy \mathfrak{h}^* .

Proof. Suppose S contains an A-set: by Theorem 4.5, it is approachable with strategy \mathfrak{g}^* .

Now suppose S does not contain an A-set; by Theorem 4.11, $S_\infty \subseteq S$ not an A-set means $T(S) < \infty$ and Proposition 4.14 can be applied. But minimax structure, combined with Proposition 3.1, grants that every halfspace-forcing counterexample (ϕ, ψ, H) has H^c 1-forcible by \mathcal{V} . Thus the stronger guarantee of Proposition 4.14 holds, and \mathfrak{h}^* is an avoidance strategy. \square

Note that without minimax structure, there exists sets which are neither avoidable nor approachable. It suffices to consider a game in \mathbb{R}^1 with $a := \inf_x \sup_y f(x, y) > \sup_y \inf_x f(x, y) =: b$, just as in Section 3. In particular, the set $H := (-\infty, (a+b)/2]$ is neither approachable nor avoidable.

5. STOCHASTIC GAMES.

The final missing piece is stochasticity; throughout this section, take \mathcal{X} and \mathcal{Y} to be families of distributions. For convenience, given a pair of distributions $(\mu, \nu) \in \mathcal{X} \times \mathcal{Y}$, let $\mathbf{E}f$ denote the map

$$(\mu, \nu) \mapsto \mathbf{E}_{X \sim \mu, Y \sim \nu}(f(X, Y)).$$

This setting is nearly an analog of the deterministic game: each player has access to the history of all distributions chosen so far, but additionally the sampled payoff $f(X_t, Y_t)$ where $(X_t, Y_t) \sim (\mu_t, \nu_t)$. Accordingly, strategies are now families of maps from these augmented histories to members of \mathcal{X} and \mathcal{Y} . There is some disagreement between authors as to the exact contents of the history; the history here was also used by Hou [6].

Definition 5.1. A set S (or, for clarity, a tuple $(\mathcal{X}, \mathcal{Y}, f, S)$) is *stochastic approachable* if

$$\exists \mathbf{g} \cdot \forall \epsilon > 0 \cdot \exists T \cdot \forall \mathbf{h} \cdot \mathbf{P}(\exists t \geq T \cdot \phi_t \notin S_\epsilon) \leq \epsilon,$$

where the probability \mathbf{P} is taken over the player history product distribution.

The crucial result is that the family of approachable sets is effectively the same, with and without randomness. Moreover, the method of proof reveals that the deterministic strategies \mathbf{g}^* and \mathbf{h}^* can be simply adapted to handle the stochastic case.

Theorem 5.2. Let \mathcal{X} and \mathcal{Y} be families of distributions over a pair of sets \mathbb{X} and \mathbb{Y} , and $f : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}^d$ and $\mathbf{E}f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$ are as above. For any set $S \subseteq \mathbb{R}^d$, $(\mathcal{X}, \mathcal{Y}, f, S)$ is stochastic approachable iff $(\mathcal{X}, \mathcal{Y}, \mathbf{E}f, S)$ is approachable.

When adapting \mathbf{g}^* and \mathbf{h}^* to the stochastic case, the following concentration result will be used. Although the independence statement may seem too strenuous, it will suffice because the adapted strategies will never care about the sampled values $X_t \sim \mu_t$ and $Y_t \sim \nu_t$, but rather only use the source distributions μ_t and ν_t .

Lemma 5.3. Let any $\epsilon > 0$ and any sequence $(Z_i)_{i=1}^\infty$ of d -dimensional independent random variables be given with $\|Z_i\| \leq \gamma$ almost surely. Define $S_n := n^{-1} \sum_{i=1}^n Z_i$. Then for any $N \geq \frac{\gamma^2}{2\epsilon^2} \ln \left(\frac{2d}{\epsilon(1 - \exp(-\epsilon^2/(2\gamma^2)))} \right)$,

$$\mathbf{P}(\exists n \geq N \cdot \|S_n - \mathbf{E}(S_n)\| \geq \epsilon) \leq \epsilon.$$

Proof. First consider any fixed $n \geq N$. By norm equivalence, $\|S_n - \mathbf{E}(S_n)\| \leq \|S_n - \mathbf{E}(S_n)\|_1$. Applying Hoeffding's inequality to any fixed coordinate j (which can vary by at most 2γ),

$$\mathbf{P}(|(S_n)_j - (\mathbf{E}(S_n))_j| \geq \epsilon) \leq 2 \exp(-2n\epsilon^2/(2\gamma)^2).$$

Unioning these events for all coordinates and all $n \geq N$, it follows that

$$\begin{aligned} \mathbf{P}(\exists n \geq N \cdot \|S_n - \mathbf{E}(S_n)\| \geq \epsilon) &\leq \sum_{n \geq N} \mathbf{P}(\|S_n - \mathbf{E}(S_n)\|_1 \geq \epsilon) \\ &\leq 2d \exp(-2N\epsilon^2/\gamma^2) \sum_{i \geq 0} \exp(-i\epsilon^2/(2\gamma^2)) \\ &\leq \epsilon(1 - \exp(-2\epsilon^2/\gamma^2)) \sum_{i \geq 0} (\exp(-\epsilon^2/(2\gamma^2)))^i. \end{aligned}$$

□

Theorem 5.2. (\Leftarrow) Suppose $(\mathcal{X}, \mathcal{Y}, \mathbf{E}f, S)$ is approachable, and let $\epsilon > 0$ be given. Define a stochastic approach strategy \mathbf{g} as follows. In every iteration, \mathbf{g} invokes \mathbf{g}^* (with schedule $\tau_t = \gamma 2^{-t}$). Since S is approachable, by Theorem 4.15 it is an A-set; thus, for $t \geq T_1 := \lceil 12\gamma^2/\epsilon^2 \rceil$, $\mathbf{E}(\phi_t) \in S_{\epsilon/2}$. Next, even though there may be dependence between X_t and Y_t across iterations, they are independent

given their distributions μ_t and ν_t . As such, Lemma 5.3 may be applied, and there exists T_2 so that $\mathbb{P}(\exists t \geq T_2 \cdot \|\phi_t - \mathbb{E}(\phi_t)\| \geq \epsilon/2) \leq \epsilon/2$. Taking $T := \max\{T_1, T_2\}$,

$$\mathbb{P}(\exists t \geq T \cdot \phi_t \notin S_\epsilon) \leq \mathbb{P}(\exists t \geq T \cdot \phi_t \notin B(\mathbb{E}(\phi_t), \epsilon/2)) + \mathbb{P}(\exists t \geq T \cdot \mathbb{E}(\phi_t) \notin S_{\epsilon/2}) \leq \epsilon.$$

(\implies) This direction is established via contrapositive: suppose $(\mathcal{X}, \mathcal{Y}, \mathbb{E}f, S)$ is not approachable, and construct an opponent strategy \mathfrak{h} by once again feeding the distribution history to a deterministic strategy, this time \mathfrak{h}^* . By Proposition 4.14, there exist $T(S)$ and $\delta' > 0$ so that $\mathbb{E}(\phi_t) \notin S_{\delta'}$ for infinitely many t ; set $\delta := \min\{\delta', 1\}$. Invoking Lemma 5.3 (with the same independence considerations as for the converse), there exists T' such that $\mathbb{P}(\forall t \geq T' \cdot \phi_t \in B(\mathbb{E}(\phi_t), \delta/2)) \geq 1 - \delta/2$, and thus, for any $T \geq \max\{T', T(S)\}$,

$$\mathbb{P}(\exists t \geq T \cdot \phi_t \notin S_{\delta/2}) \geq 1 - \delta/2.$$

□

Remark 5.4. Note $\mathbb{E}f$ is bilinear, and suppose \mathcal{X} and \mathcal{Y} are compact convex. This grants minimax structure (cf. Appendix D), and one may view the stochasticity as an operator embedding a tricky game into a more structured setting.

APPENDIX A. AN EXAMPLE.

Many beautiful approachability examples can be found elsewhere. Here are two favorites:

- Blackwell [1] presents a game which is neither approachable nor excludable;
- Spinat [10] presents an approachable nonconvex set containing no smaller approachable set.

This section presents a very simple game to demonstrate the relationship of 1-forcing, 2-forcing, and approachability. For its duration, fix

$$\begin{aligned} \mathcal{X} &:= \mathcal{Y} := \{[1-\alpha] : \alpha \in [0, 1]\}, \\ f(x, y) &:= \begin{bmatrix} x_1 y_1 \\ x_2 y_2 \end{bmatrix}. \end{aligned}$$

Since this game is symmetric, it suffices to consider the perspective of one player. Notice that the minimal 1-forcible sets are precisely

$$L_x := f(x, \mathcal{Y}) = \left\{ \begin{bmatrix} \alpha x_1 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ (1-\alpha)x_2 \end{bmatrix} : \alpha \in [0, 1] \right\};$$

these are minimal in the sense that every 1-forcible set must contain L_x for some x , and each L_x has no 1-forcible proper subsets.

For a convex approachable set that is not 1-forcible, consider

$$S_0 := \{(\alpha, \alpha) : \alpha \in [0, 1/2]\}.$$

Since every halfspace containing S_0 also contains some L_x , Theorem 4.15 grants approachability.

This game is a tensor of order 3, and the results of Appendix D grant it has the minimax property. But there exist 2-forcible sets which are not halfspaces, thus not covered by Proposition 3.1, and are not approachable. Consider in particular the set

$$S_1 := ([1/2, 1] \times \{0\}) \cup (\{0\} \times (1/2, 1]).$$

Since S_1 intersects every set L_x , it follows that it is 2-forcible; by the aforementioned symmetry of f , it is 2-forcible for \mathcal{X} and for \mathcal{Y} . On the other hand since S_0 is approachable but $\rho(S_0, S_1) > 0$, the \mathcal{Y} -player can play \mathfrak{g}^* to approach S_0 , meaning S_1 is not approachable.

Finally, as per Remark 4.7, there exist tricky supersets of A-sets for which it is not clear how to construct an approach strategy (without resorting to determining an inner A-set, and invoking \mathfrak{g}^* on it). In particular, choose any $x_1, x_2 \in \mathcal{X}$ with $x_1 \neq x_2$, and construct S_2 by removing a small piece from L_{x_2} in $L_{x_1} \cup L_{x_2}$. Any approach strategy for S_2 must know to avoid L_{x_2} ; otherwise, the \mathcal{Y} -player could let arbitrarily many ϕ_t stay inside the disconnected piece of L_{x_2} , and then begin forcing points inside L_{x_1} , thus taking ϕ_t outside of S_2 . Since approachability requires a uniform bound over




FIGURE 4. Two cases in the proof of Lemma B.1.

all opponent strategies, this means the provided strategy for \mathcal{X} is not an approach strategy. Thus, it is necessary for the \mathcal{X} -player strategy to knowingly avoid this disconnected piece of L_{x_2} .

APPENDIX B. GEOMETRIC FACTS.

This section collects a few geometric facts requiring careful proof.

B.1. Proof of Lemma 4.8. Throughout this section, suppose (ϕ, ψ, H) is merely a halfspace-forcing candidate; 1-forcing of H by \mathcal{X} or \mathcal{Y} will only come into play at the end. As in Lemma 4.8, set $\tau := \rho(\psi, H^c) > 0$, and set $H' := \overline{H^c}$.

Lemma B.1. *Set $\epsilon_0 := \tau(1 - \gamma/\sqrt{\gamma^2 + \tau^2})$, and $\epsilon_1 := \tau\epsilon_0/(2\sqrt{4\gamma^2 + \tau^2})$. For any $z \in H' \cap f(\mathcal{X}, \mathcal{Y})$, there exists $\eta \in (0, 1)$ such that $B(\eta\psi + (1 - \eta)z, \epsilon_1) \subseteq B(\phi, \rho(\phi, \psi))$ and $\rho(\eta\psi + (1 - \eta)z, H') \geq \epsilon_1$.*

Proof. First, it suffices to consider z on the boundary of H' : given $z' \in (H')^o$ with z denoting the intersection of $[\psi, z']$ with the boundary of H , the desired η for z can be converted into $\eta' \in (0, 1)$ for z' satisfying $\eta'\psi + (1 - \eta')z' = \eta\psi + (1 - \eta)z$.

Second, discard the scenario that $\{\psi, z, \phi\}$ are collinear: since $\epsilon_1 < \tau/2 \leq \rho(\phi, \psi)/2$, taking η close to $1/2$ suffices.

Third, it suffices to exhibit an $\eta \in [0, 1)$ satisfying the single property $B(\eta\psi + (1 - \eta)z, \epsilon_0) \subseteq B(\phi, \rho(\phi, \psi))$. This η may potentially violate the second condition above: perhaps $\rho(\eta\psi + (1 - \eta)z, H') < \epsilon_1$. To see how this can be adjusted, set $z' := \eta\psi + (1 - \eta)z$ and ϕ' to be the projection of z' onto $[\phi, \psi]$; thus the right triangle $\{\psi, \phi', z'\}$ has short sides with lengths $\rho(\psi, \phi') \geq \tau - \epsilon_1 \geq \tau/2$ and $\rho(\phi', z') \leq \gamma$. By similarity of triangles, the point z'' along $[z', \psi]$ which is $\epsilon_0/2$ away from z' must satisfy

$$\begin{aligned} \rho(z'', H^c) &\geq \rho(z'', [z', \phi']) = \rho(z'', z') \left(\frac{\rho(\phi', \psi)}{\rho(z', \psi)} \right) = \frac{\epsilon_0}{2} \left(\frac{1}{\sqrt{1 + (\rho(\phi', z')/\rho(\phi', \psi))^2}} \right) \\ &\geq \frac{\epsilon_0}{2} \left(\frac{\tau}{\sqrt{\tau^2 + 4\gamma^2}} \right) = \epsilon_1. \end{aligned}$$

Meanwhile, $B(z'', \epsilon_1) \subseteq B(z'', \epsilon_0/2) \subseteq B(z', \epsilon_0)$, meaning both properties are satisfied. And since $z'' \in (\psi, z)$, this grants an $\eta \in (0, 1)$ with all desired properties.

The remainder of the proof will be divided into the two cases, as in Figure 4, whether $\angle(\phi, z, \psi) < \pi/2$ or not. Set ϕ^* to be the intersection of $[\phi, \psi]$ with the boundary of H , meaning $\rho(\psi, \phi^*) = \tau$.

Suppose $\angle(\phi, z, \psi) < \pi/2$, and designate z^* as the point along $[z, \psi]$ satisfying $[z^*, \psi] \perp [z^*, \phi]$, as in Figure 4a. This z^* will provide the eventual η . By similarity of triangles $\{\psi, z^*, \phi\}$ and $\{\psi, \phi^*, z\}$,

$$\frac{\rho(z^*, \phi)}{\rho(\psi, \phi)} = \frac{\rho(z, \phi^*)}{\sqrt{\rho(z, \phi^*)^2 + \rho(\psi, \phi^*)^2}}.$$

Rearranging and making use of $\rho(\phi^*, z) \leq \gamma$ and $\rho(\psi, \phi) \geq \rho(\psi, \phi^*) = \tau$, the desired inequality is

$$\begin{aligned} \rho(\phi, \psi) - \rho(z^*, \phi) &= \rho(\phi, \psi) \left(1 - \frac{1}{\sqrt{1 + \rho(\psi, \phi^*)^2 / \rho(z, \phi^*)^2}} \right) \geq \rho(\phi, \psi) \left(1 - \frac{1}{\sqrt{1 + \tau^2 / \gamma^2}} \right) \\ &\geq \tau \left(1 - \frac{\gamma}{\sqrt{\gamma^2 + \tau^2}} \right) = \epsilon_0. \end{aligned}$$

Finally, consider $\angle(\phi, z, \psi) \geq \pi/2$. This time, place z^* along the line through $\{z, \phi^*\}$ so that $[z^*, \phi] \perp [z^*, \psi]$ as in Figure 4b. By construction, $\rho(z^*, \phi) \geq \rho(z, \phi)$, thus it suffices to upper bound $\rho(z^*, \phi)$, and z will be the desired point, meaning $\eta = 0$. By similar of triangles $\{\psi, \phi, z^*\}$ and $\{z^*, \phi, \phi^*\}$,

$$\frac{\rho(\psi, \phi)}{\rho(\phi, z^*)} = \frac{\rho(z^*, \phi)}{\rho(\phi, \phi^*)}.$$

Using $\rho(\psi, \phi) \geq \rho(\psi, \phi^*) = \tau$ and $\rho(\phi^*, \phi) \leq \gamma$,

$$\rho(z^*, \phi) = \rho(\psi, \phi) \sqrt{\frac{\rho(\phi, \phi^*)}{\rho(\phi, \phi^*) + \rho(\phi^*, \psi)}} = \rho(\psi, \phi) \sqrt{\frac{1}{1 + \rho(\phi^*, \psi) / \rho(\phi, \phi^*)}} \leq \rho(\psi, \phi) \sqrt{\frac{1}{1 + (\tau/\gamma)^2}},$$

with the remainder as before. \square

Lemma B.2. Set $\epsilon_2 := \epsilon_1/2$. For any $p \in B(\psi, \epsilon_2)$, and any $z \in H' \cap f(\mathcal{X}, \mathcal{Y})$, there exists $\eta \in (0, 1)$ such that $B(\eta p + (1 - \eta)z, \epsilon_2) \subseteq B(\phi, \rho(\phi, \psi))$ and $\rho(\eta p + (1 - \eta)z, H') \geq \epsilon_2$.

Proof. Choosing the η granted by Lemma B.1,

$$\rho(\eta p + (1 - \eta)z, \phi) = \|\eta(p - \psi + \psi) + (1 - \eta)z - \phi\| \leq \underbrace{\eta \|p - \psi\|}_{\leq \epsilon_2} + \underbrace{\|\eta\psi + (1 - \eta)z - \phi\|}_{\leq \rho(\phi, \psi) - \epsilon_1} \leq \rho(\phi, \psi) - \epsilon_2.$$

For the other property,

$$\rho(\eta p + (1 - \eta)z, H') = \inf_{q \in H'} \|\eta p + (1 - \eta)z - q\| + \|\eta(\psi - p)\| - \|\eta(\psi - p)\| \geq \rho(\eta\psi + (1 - \eta)z, H') - \epsilon_2.$$

\square

Lemma B.3. Set $\epsilon_3 := \epsilon_2/2 = \epsilon$. Let any $p \in B(\psi, \epsilon_2)$, $T \geq \lceil 8\gamma/\epsilon_3 \rceil$, and $(z_i)_{i=1}^N \in (H' \cap f(\mathcal{X}, \mathcal{Y}))^N$ with $N = \lceil T\gamma\epsilon_3/8 \rceil$ be given. Then there exists an integer $M \leq N$ such that

$$\rho\left(\frac{Tp + \sum_{i=1}^M z_i}{T + M}, \phi\right) \leq \rho(\phi, \psi) - \epsilon_3.$$

Proof. For every $z \in H' \cap f(\mathcal{X}, \mathcal{Y})$, let $\eta_z \in (0, 1)$ be the value granted by Lemma B.2 so that

$$B(\eta_z p + (1 - \eta_z)z, \epsilon_2) \subseteq B(\phi, \rho(\phi, \psi)) \cap H.$$

Next, set $U_z := B(\eta_z p + (1 - \eta_z)z, \epsilon_3)$, and \mathcal{U} to be the convex hull of the union of all U_z ; it follows that $\mathcal{U} \subseteq B(\phi, \rho(\phi, \psi) - \epsilon_2)$ and $\rho(\mathcal{U}, H^c) = \inf_{u \in \mathcal{U}} \rho(u, H^c) \geq \epsilon_2$.

For every $i \leq N$, consider the partial averages

$$q_i := \frac{Tp + \sum_{j=1}^i z_j}{T + i},$$

and take w_i to denote the point on the boundary of H' so that $q_i \in [p, w_i]$ (or $w_i = q_i$ if $q_i \in H'$). Correspondingly, choose η_i as provided by Lemma B.2 so that $\eta_i p + (1 - \eta_i)w_i \in \mathcal{U}$, and set $\mu_i \in [0, 1]$ so that $q_i = \mu_i p + (1 - \mu_i)w_i$. Since $N \geq T\gamma\epsilon_3/8$ and $\gamma > 0$,

$$\left\| q_N - \frac{1}{N} \sum_{i=1}^N z_i \right\| = \frac{1}{T + N} \left\| Tp + \sum_{i=1}^N z_i - \frac{T + N}{N} \sum_{i=1}^N z_i \right\| = \frac{T}{T + N} \left\| p - \frac{1}{N} \sum_{i=1}^N z_i \right\| \leq \frac{T\gamma}{T + N} \leq \epsilon_3/8,$$




FIGURE 5. Objects in the proof of Lemma B.5.

meaning in particular that $\rho(q_N, H') \leq \epsilon_3/8$, so has gotten beyond \mathcal{U} , thus $\mu_N \geq \eta_N$. The final step will be to show $(q_i)_{i=1}^N$ must have actually passed through \mathcal{U} .

Now let k be the first index such that $\mu_k \geq \eta_k$, meaning $\mu_{k-1} < \eta_{k-1}$. Suppose contradictorily that $q_k \notin \mathcal{U}$ and $q_{k-1} \notin \mathcal{U}$. Since $T \geq 8\gamma/\epsilon_3$, it follows that $\|q_k - q_{k-1}\| \leq \epsilon_3/8$. This implies that the ball centered at $\eta_i p + (1 - \eta_i)w_k$ with radius $\epsilon_3/4$ (which is within $B(\eta_i p + (1 - \eta_i)w_k, \epsilon_3) \subseteq \mathcal{U}$) contains a point along the line $[p, q_{k-1}]$, and as such $q_{k-1} \in \mathcal{U}$, a contradiction. Thus the desired M exists. \square

Lemma 4.8. It is given that (ϕ, ψ, H) is a halfspace-forcing example, by which it follows that H^c and hence H' can be 2-forced by \mathcal{Y} . Given a sequence $(x_i)_{i=1}^N$, this grants $(y_i)_{i=1}^N$ with $f(x_i, y_i) \in H' \cap f(\mathcal{X}, \mathcal{Y})$, whereby Lemma B.3 may be applied, and the result follows.

Now suppose the stronger condition that H^c (and H') can be 1-forced by \mathcal{Y} . Thus there exists a single \bar{y} so that the choice $y_i = \bar{y}$ grants $f(x_i, y_i) \in H' \cap f(\mathcal{X}, \mathcal{Y})$, whatever the choice of $x_i \in \mathcal{X}$. Once again, Lemma B.3 may be applied. \square

B.2. Halfspace-forcing Counterexamples of Similar Sets. This subsection will use the existence of a halfspace-forcing counterexample (ϕ, ψ, H) on a set S to produce another counterexample on some similar set S' . It will be supposed that $\rho(\phi, \psi) = \rho(\psi, H^c) =: \tau$; this comes without loss of generality, since otherwise $\rho(\phi, \psi) > \tau$, in which case some other ϕ' may be chosen along the segment $[\phi, \psi]$, but closer to ψ , and still forming a halfspace-forcing counterexample (ϕ', ψ, H) .

To this end, given any pair (ϕ, q) , define $H(\phi, q)$ to be the halfspace with normal $q - \phi$ having ϕ on its boundary; that is,

$$(B.4) \quad H(\phi, q) := \{z \in \mathbb{R}^d : \langle z, q - \phi \rangle \leq \langle \phi, q - \phi \rangle\}.$$

For instance, after the adjustment placing ϕ on the boundary of H , $H = \overline{H(\phi, \psi)^c}$, and $(\phi, \psi, \overline{H(\phi, \psi)^c})$ is a halfspace-forcing counterexample.

For any pair (ϕ, q) and $\delta \in (0, \rho(\phi, q))$, let $R_\delta(\phi, q)$ be a shell of radius $\rho(\phi, q)$ and width δ around ϕ :

$$R_\delta(\phi, q) := B(\phi, \rho(\phi, q)) \setminus B(\phi, \rho(\phi, q) - \delta)^o = \{z \in \mathbb{R}^d : \rho(\phi, q) \geq \rho(\phi, z) \geq \rho(\phi, q) - \delta\}.$$

Lemma B.5. *Let ϕ, ψ be given where $H(\phi, \psi)$ can be 2-forced by \mathcal{Y} , and set $\tau := \rho(\phi, \psi)$ and $\phi' := (\phi + \psi)/2$. Then there exists $\delta > 0$ so that every $q \in B(\phi', \tau/2) \cap R_\delta$ satisfies*

- (1) $\rho(q, \psi) \leq \tau/4$,
- (2) $H(\phi', q)$ can be 2-forced by \mathcal{Y} .

Many of the relevant quantities appear in Figure 5.

Proof. For any $\delta \in (0, \tau/2)$, consider two maximization problems, where capitalization refers to the presence of multiple optima:

$$\begin{aligned} q_\delta &\in \operatorname{Argmax}\{\angle(\psi, \phi, q) : q \in B(\phi', \tau/2) \cap R_\delta\}, \\ q'_\delta &\in \operatorname{Argmax}\{\rho(\psi, q) : q \in B(\phi', \tau/2) \cap R_\delta\}. \end{aligned}$$

By rotational symmetry of all quantities around $[\phi, \psi]$, it suffices to consider q_δ, q'_δ which lie in the same plane, and lie on or above $[\phi, \psi]$. In this case, $q_\delta = q'_\delta$: this follows since the farthest point from ψ must lie on the inner edge of R_δ , and it must be as high above $[\phi, \psi]$ as possible.

As such, to establish the first desired property, note that $\rho(\psi, q_\delta)$ decreases continuously as $\delta \downarrow 0$, and thus there exists δ_1 so that $\rho(\psi, q_{\delta_1}) \leq \tau/4$. Since q_δ is the farthest point, the property follows.

For the second property, since H^c can be 2-forced by \mathcal{Y} , it follows that $H^c \cap f(\mathcal{X}, \mathcal{Y})$ can be 2-forced by \mathcal{Y} . Now consider any point $q \in B(\phi, \tau/2) \cap R_\delta$, and consider the halfspace $H(\phi', q)$ as defined above. Once again, it suffices by rotational symmetry of the relevant quantities around $[\phi, \psi]$ to consider q as lying in a plane above $[\psi, \phi']$ and $H(\phi', q)$ appearing as a line within this plane. If $H(\phi', q)$ does not cross $H^c \cap f(\mathcal{X}, \mathcal{Y})$ within this plane, then $H(\phi', q)$ can be 2-forced by \mathcal{Y} .

To this end, since $\angle(\psi, \phi', q_\delta) = \max\{\angle(\psi, \phi, q) : q \in B(\phi', \tau/2) \cap R_\delta\}$, it suffice to prove that $H(\phi', q_\delta)$ can be 2-forced. But as $\delta \downarrow 0$, the preceding analysis grants q_δ becomes arbitrarily close to ψ , thus the normal of $q_\delta - \phi'$ of $H(\phi', q)$ becomes increasingly close to that of $H(\phi, \psi)$, whereas they intersect $[\phi, \psi]$ at points $\tau/2$ apart; so there must exist a δ_2 sufficiently small.

Setting $\delta := \min\{\delta_1, \delta_2\}$, the result follows. \square

Proposition B.6. *Let $S \subseteq \mathbb{R}^d$ be given with halfspace-forcing counterexample (ϕ, ψ, H) , and set $\tau := \rho(\phi, \psi)$. Then there exists $\delta > 0$ so that every $S' \subseteq \mathbb{R}^d$ satisfying $\Delta(S, S') < \delta$ has a halfspace-forcing counterexample $(\phi'_s, q', \overline{H(\phi'_s, q')^c})$ with $\rho(H(\phi'_s, q'), q') \geq \tau/4$, where $H(\cdot, \cdot)$ is as in Equation (B.4).*

Proof. Let $\delta_0 > 0$ be as provided by Lemma B.5, set $\delta := \delta_0/2$, and let S' be any set satisfying $\Delta(S, S') < \delta_0/2$. Also following Lemma B.5, set $\phi' := (\psi + \phi)/2$. Define

$$\begin{aligned} \phi_s &:= \phi - \frac{\delta_0}{2} \left(\frac{\phi - \psi}{\|\phi - \psi\|} \right), \\ \phi'_s &:= \phi' - \frac{\delta_0}{2} \left(\frac{\phi - \psi}{\|\phi - \psi\|} \right), \\ \psi_s &:= \psi - \frac{\delta_0}{2} \left(\frac{\phi - \psi}{\|\phi - \psi\|} \right); \end{aligned}$$

these are just copies of ϕ, ϕ', ψ shifted by $\delta_0/2$ along the direction to ψ from ϕ . As such, there is a bijection (via this shift) granting

$$B(\phi', \tau/2) \cap R_{\delta_0}(\phi, \psi) \ni q \mapsto q_s \in B(\phi'_s, \tau/2) \cap R_{\delta_0}(\phi_s, \psi_s).$$

It follows from Lemma B.5 that every such q_s satisfies $q_s \in B(\psi_s, \tau/4)$, and $H(\phi'_s, q_s)$, which contains $H(\phi', q)$, is 2-forcible by \mathcal{Y} .

Since $B(\phi, \tau)^o \cap S = \emptyset$, the triangle inequality grants

$$\begin{aligned} B(\phi_s, \tau - \delta_0/2)^o \cap S &= \emptyset, \\ B(\phi_s, \tau - \delta_0)^o \cap S' &= \emptyset. \end{aligned}$$

Combining this with the definition of $R_{\delta_0}(\phi_s, \psi_s)$,

$$\begin{aligned} S \cap B(\phi_s, \tau) &\subseteq R_{\delta_0}(\phi_s, \psi_s), \\ S' \cap B(\phi_s, \tau) &\subseteq R_{\delta_0}(\phi_s, \psi_s). \end{aligned}$$

So by the choice of δ_0 and guarantees of Lemma B.5 (points q_s in this cap satisfy $q_s \in B(\psi_s, \tau/4)$),

$$\begin{aligned} S \cap B(\phi'_s, \tau/2) &\subseteq B(\psi_s, \tau/4) \cap B(\phi'_s, \tau/2) \cap R_{\delta_0}(\phi_s, \psi_s) \subseteq B(\psi_s, \tau/4) \cap R_{\delta_0}(\phi_s, \psi_s), \\ S' \cap B(\phi'_s, \tau/2) &\subseteq B(\psi_s, \tau/4) \cap B(\phi'_s, \tau/2) \cap R_{\delta_0}(\phi_s, \psi_s) \subseteq B(\psi_s, \tau/4) \cap R_{\delta_0}(\phi_s, \psi_s). \end{aligned}$$

In particular, if $S' \cap B(\phi'_s, \tau/2)$ is nonempty, then the closest point $q' \in S'$ to ϕ'_s will fall within $B(\psi_s, \tau/4) \cap R_{\delta_0}(\phi_s, \psi_s)$, and satisfies all desired properties: $H(\phi_s, q')$ was shown earlier to be 2-forced by \mathcal{Y} , thus $(\phi'_s, q', \overline{H(\phi'_s, q')})$ is a halfspace-forcing counterexample for S' ; combining $\rho(q', \psi_s) \leq \tau/4$ with $\rho(\phi'_s, \psi_s) = \tau/2$ grants $\rho(H(\phi'_s, q'), q') \geq \tau/4$.

To this end, note that by construction that

$$B(\psi, \delta_0/2) \subseteq R_{\delta_0}(\phi_s, \psi_s).$$

S' satisfies $\Delta(S, S') < \delta_0/2$, meaning there exists q' inside $S' \cap R_{\delta_0}(\phi_s, \psi_s)$. The result follows. \square

APPENDIX C. A LIMIT PROPERTY.

The proof of Spinat's [10] Lemma 1 appears to be incomplete. The statement is interesting, so this appendix establishes a strengthened form, albeit by different means.

Lemma C.1. *Let a sequence of compact approachable sets $(U_i)_{i \geq 1}$ be given. If $(U_i)_{i \geq 1}$ is convergent in Hausdorff metric, then its limit U is a compact approachable set.*

Proof. Since each U_i is approachable, Theorem 4.15 guarantees that it contains a compact (nonempty) A-set A_i . Completeness of the Hausdorff metric on compact nonempty sets grants (perhaps by passing to a subsequence) the $(A_i)_{i \geq 1}$ have a limiting (compact nonempty) set A .

It must be the case that $A \subseteq U$, since otherwise by compactness there exists $z \in A \setminus U$ with some $\rho(z, U) =: \delta > 0$. Now consider j (with respect to the above subsequence) large enough for $\Delta(U_j, U) \leq \delta/4$ and $\Delta(A_j, A) \leq \delta/2$; it must be the case that $A_j \not\subseteq U_j$, a contradiction.

Finally, note that A is an A-set. Suppose contradictorily that (ϕ, ψ, H) is a counterexample to halfspace-forcibility with $\tau := \rho(\psi, H^c)$. By Proposition B.6, there must exist a tiny $\delta > 0$ so that every A' with $\Delta(A, A') < \delta$ has halfspace-forcing counterexamples, whereas there exists an A-set A_j with $\Delta(A, A_j) < \delta$.

Since U contains an A-set, Theorem 4.15 grants that it is approachable. \square

APPENDIX D. FUNCTION FAMILIES SATISFYING THE MINIMAX PROPERTY.

Recall Sion's [9] minimax theorem, as stated by Komiya [7], with a minor tightening to compact \mathcal{Y} for applicability here.

Theorem D.1 (Sion [9]). *Let $g : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ be given where \mathcal{X} and \mathcal{Y} are compact convex subsets of linear topological spaces, $g(\cdot, y)$ is quasiconvex and lower semi-continuous for every $y \in \mathcal{Y}$, and $g(x, \cdot)$ is quasiconcave and upper semi-continuous for every $x \in \mathcal{X}$. Then*

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} g(x, y) = \max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} g(x, y);$$

in particular, each optimization is attainable.

Notice that Sion's Theorem may be applied to the case that \mathcal{X} and \mathcal{Y} are families of distributions.

Proposition D.2. *Let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$ be given. $\langle f(\cdot, y), \lambda \rangle$ is quasiconvex and lower semi-continuous for every $(y, \lambda) \in \mathcal{Y} \times \mathbb{R}^d$ and $\langle f(x, \cdot), \lambda \rangle$ is quasiconcave and upper semi-continuous for every $(x, \lambda) \in \mathcal{X} \times \mathbb{R}^d$ iff $\langle f(\cdot, \cdot), \lambda \rangle$ is continuous and monotonic in each parameter.*

Proof. (\implies) Fix y and λ ; since both $\langle f(\cdot, y), \lambda \rangle$ and $\langle f(\cdot, y), -\lambda \rangle$ are quasiconvex and lower semi-continuous, it follows that each is continuous, and unrolling quasiconvexity grants, for any $\alpha \in [0, 1]$,

$$\min\{\langle f(x_1, y), \lambda \rangle, \langle f(x_2, y), \lambda \rangle\} \leq \langle f(\alpha x_1 + (1 - \alpha)x_2, y), \lambda \rangle \leq \max\{\langle f(x_1, y), \lambda \rangle, \langle f(x_2, y), \lambda \rangle\},$$

which is the statement of monotonicity. An analogous argument holds for every x and λ .

(\impliedby) Continuity implies upper and lower semi-continuity, and every monotonic function g satisfies, for every x_1, x_2 and $\alpha \in [0, 1]$, $\min\{g(x_1), g(x_2)\} \leq g(\alpha x_1 + (1 - \alpha)x_2) \leq \max\{g(x_1), g(x_2)\}$. \square

In the stricter scenario of convexity/concavity, the resulting function family is vastly more constrained.

Proposition D.3. *Let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$ be given. $\langle f(\cdot, y), \lambda \rangle$ is convex for every $(y, \lambda) \in \mathcal{Y} \times \mathbb{R}^d$ and $\langle f(x, \cdot), \lambda \rangle$ is concave for every $(x, \lambda) \in \mathcal{X} \times \mathbb{R}^d$ iff $f(\cdot, \cdot)$ is affine in each parameter.*

Proof. (\implies) Since for any λ and y , both $\langle f(\cdot, y), \lambda \rangle$ and $\langle f(\cdot, y), -\lambda \rangle$ are convex, it follows that $\langle f(\cdot, y), \lambda \rangle$ is affine. Thus let any $\alpha \in \mathbb{R}$ and $x_1, x_2 \in \mathcal{X}$ be given; for any $y \in \mathcal{Y}$,

$$\begin{aligned} f(\alpha x_1 + (1 - \alpha)x_2, y) &= \sum_{i=1}^d \langle f(\alpha x_1 + (1 - \alpha)x_2, y), \mathbf{e}_i \rangle \mathbf{e}_i \\ &= \sum_{i=1}^d (\alpha \langle f(x_1, y), \mathbf{e}_i \rangle + (1 - \alpha) \langle f(x_2, y), \mathbf{e}_i \rangle) \mathbf{e}_i \\ &= \alpha f(x_1, y) + (1 - \alpha)f(x_2, y). \end{aligned}$$

Repeating this proof from the perspective of $f(x, \cdot)$, the result follows.

(\impliedby) Let $\alpha \in [0, 1]$ and $x_1, x_2 \in \mathcal{X}$ be given. Then, for any $\lambda \in \mathbb{R}^d$,

$$\langle f(\alpha x_1 + (1 - \alpha)x_2, y), \lambda \rangle = \langle \alpha f(x_1, y) + (1 - \alpha)f(x_2, y), \lambda \rangle = \alpha \langle f(x_1, y), \lambda \rangle + (1 - \alpha) \langle f(x_2, y), \lambda \rangle.$$

Again, the proof for \mathcal{Y} is analogous. \square

Acknowledgement. The author thanks his advisor, Sanjoy Dasgupta, for discussions and support.

REFERENCES

1. David Blackwell, *An analog of the minimax theorem for vector payoffs*, Pacific Journal of Mathematics **6** (1956), no. 1, 1–8.
2. ———, *Controlled random walks*, invited address, Institute of Mathematical Statistics Meeting, Seattle, Washington, 1956.
3. Jonathan Borwein and Adrian Lewis, *Convex analysis and nonlinear optimization*, Springer Publishing Company, Incorporated, 2000.
4. Dean P. Foster and Rakesh V. Vohra, *Asymptotic calibration*, Biometrika **85** (1998), no. 2, 379–390.
5. Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal, *Fundamentals of convex analysis*, Springer Publishing Company, Incorporated, 2001.
6. Tien-Fang Hou, *Approachability in a two-person game*, The Annals of Mathematical Statistics **42** (1971), no. 2, 735–744.
7. Hidetoshi Komiya, *Elementary proof for sion's minimax theorem*, Kodai Mathematical Journal **11** (1988), no. 1, 5–7.
8. James R. Munkres, *Topology*, 2nd ed., Prentice Hall, 2000.
9. Maurice Sion, *On general minimax theorems*, Pacific Journal of Mathematics **8** (1958), no. 1, 171–176.
10. Xavier Spinat, *A necessary and sufficient condition for approachability*, Mathematics of Operations Research **27** (2002), no. 1, 31–44.
11. Nicolas Vieille, *Weak approachability*, Mathematics of Operations Research **17** (1992), no. 4, 781–791.

E-mail address: mtelgars@cs.ucsd.edu